




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04L 29/06, G06F 1/00		A1	(11) International Publication Number: WO 99/41888
			(43) International Publication Date: 19 August 1999 (19.08.99)
(21) International Application Number: PCT/US99/03382 (22) International Filing Date: 17 February 1999 (17.02.99) (30) Priority Data: 09/024,576 17 February 1998 (17.02.98) US (71) Applicant: SECURE COMPUTING CORPORATION [US/US]; 2675 Long Lake Road, Roseville, MN 55113 (US). (72) Inventors: VIETS, Richard, R.; 1360 Nottingham Drive, Naples, FL 34109 (US). MOTES, David, G.; 10301 Westside Lane, Bonita Springs, FL 34135 (US). GREVE, Paula, Budig; Apartment 205, 2516 Silver Lane, St. Anthony, MN 55421 (US). HERBERG, Wayne, W.; 10600 487th Street, Rush City, MN 55069 (US). (74) Agent: VIKSNINS, Ann, S.; Schwegman, Lundberg, Woessner & Kluth, P.O. Box 2938, Minneapolis, MN 55402 (US).		(81) Designated States: European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	

(54) Title: SYSTEM AND METHOD FOR CONTROLLING ACCESS TO STORED DOCUMENTS



(57) Abstract

A system and method of limiting access from an external network to documents stored on an internal network. A client list is built in which each client is assigned to one or more roles. Each role has access to one or more documents as defined on a document list. A request from an external network is reviewed and, if possible, the request is associated with a client on the client list. The requested document is then compared to the document list associated with the client's role and, if the requested document is in the list of documents available to a client in the client's role, the requested document is fetched, cleaned and sent to the client.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

System and Method for Controlling Access to Stored Documents

5

Background of the Invention

Field of the Invention

The present invention relates to systems and methods for controlling communication between networks, and in particular to a system and method for limiting access to documents stored on an internal network.

10

Background Information

Businesses today are acting cooperatively to achieve compatible business goals. For example, companies are using just-in-time manufacturing techniques to reduce overhead. To make this work, companies rely heavily on the ability of their suppliers to provide materials when needed.

15

At the same time, in this digital age business executives have become accustomed to receiving information from a number of sources both inside and outside the company almost instantaneously. They rely on such information to drive their day-to-day management decisions.

20

In order to provide outside organizations with relevant information in a timely manner, many companies have expanded their order-processing departments to handle increased call volumes. In this environment, outside partners call into the company's order-processing department to request specific information. This requires an employee to be available to answer calls, pull up information and verbally convey information to the partner. This option is very expensive, slow, and offers a poor level of service. What is needed is a system and method of streamlining the flow of information between partner companies while limiting access to company proprietary information.

25

The Internet provides one possible solution to this problem. The nature of the Internet makes it an ideal vehicle for organizations to communicate and share information. The Internet offers low cost universal access to information. Because of this, Internet transactions are expected to more than quadruple over the next two years, and partner communications via the Internet will almost double. Companies have begun to look to the Internet as a medium allowing

30

quick, easy and inexpensive to business partners. To date, however, their Internet options have been limited.

One solution is to give business partners access to the company internal network. Companies are hesitant to do this, however, since such access, if
5 abused, can lead to the disclosure of company sensitive information.

Another solution is to replicate necessary information to a web server located outside the company's firewall. Such an approach does allow organizations direct access to the information while at the same time limiting their access to company sensitive information. For this environment to work,
10 however, the MIS department must manually transfer information from the internal network to the external server. Therefore, while this option offers organizations direct access to necessary data, that information can be 24 to 48 hours old. When dealing with just-in-time inventory levels and large dollar amounts, 24 hours is too late. This option also creates a bottleneck in MIS,
15 redundancy of data, and decreased data integrity.

What is needed is a system and method for giving controlled access to designated documents stored on the internal network while restricting access to company sensitive information.

Summary of the Invention

20 The present invention is a system and method of limiting access from an external network to documents stored on an internal network. A client list is built in which each client is assigned to one or more roles. Each role has access to one or more documents as defined on a document list. A request from an external network is reviewed and, if possible, the request is associated with a
25 client on the client list. The requested document is then compared to the document list associated with the client's role and, if the requested document is in the list of documents available to a client in the client's role, the requested document is fetched, cleaned and sent to the client.

According to another aspect of the present invention, a document control
30 system is described. The document control system includes an internal network, an external interface, a document server connected to the internal network, and a document control server connected to the internal network and to the external

interface. The document server controls access to a plurality of documents, including a first document. The document control server includes a go list processor for determining if the user has authorization to access said first document and a document processor for reading the first document from the document server, cleaning the first document and forwarding a clean version of said first document to the user. In operation, the document control server receives a document request from the external interface for the first document, determines a user associated with the document request, authenticates the user, determines if the user has authorization to access said first document and, if authorized, reads the first document from the document server, cleans the first document and forwards a clean version of said first document to the user.

Brief Description of the Drawings

In the drawings, where like numerals refer to like components throughout the several views,

Figure 1 shows a document access system;

Figure 2 is a flow diagram illustrating operations performed by the document access system of Figure 1;

Figure 3 shows a document control server which can be used in the document access system shown in Figure 1;

Figure 4 is a document access system which includes a firewall;

Figure 5 is a document access system in which the document control server is placed in a third network;

Figure 6 is a document access system in which the document control server is placed on the external network; and

Figure 7 is an example of a tree structure representation which could be used to aid the data owner in the selection of permitted URLs.

Description of the Preferred Embodiments

In the following detailed description of the preferred embodiments, reference is made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration specific embodiments in which the invention may be practiced. It is to be understood that other embodiments may

be utilized and structural changes may be made without departing from the scope of the present invention.

As noted above, corporations today are required by customers to deliver information such as price changes, new product data, manufacturing data, and customer support electronically. Competition is driving firms to work with partners through tight connections to internal systems. Allowing access, however, to such data in an efficient, manageable, and secure manner presents challenges. Companies go to great lengths to set up order processing departments and replicate large quantities of internal data to an external Internet server. These efforts are not only inefficient, but usually result in redundancy of data, decreased network integrity, and a bottleneck in the MIS department.

The present invention solves this problem by allowing specified external users controlled, customized, and secure access to the company's intranet without complex network infrastructure modifications. Further, the present invention permits one to control the parts of a Web server that are accessible to a Business Partner with only minimal intervention by IS personnel. (The term "Business Partner" is used in the following discussion to describe an external user who needs access to data such as Web pages which are not generally available to the public, but who also should not have unlimited to a company's intranet Web services.)

A document access system for giving controlled access to designated documents stored on the internal network while restricting access to company sensitive information is shown in Fig. 1. In Fig. 1 document access system 10 includes a document control server 12, a document server 14, an external interface 16 and one or more internal workstations 18. Document control server 12, document server 14, external interface 16 and internal workstations 18 are interconnected via internal network 20. Document server 14 reads and writes documents to storage 20. Requests for documents arrive at external interface 16 and are forwarded to document control server 12 for execution. In one embodiment external interface 16 includes a router used to form an Internet connection. In another embodiment, external interface 16 includes a direct

connection interface such as formed by one or more modems used for direct dial-up by business partners wishing to access their data.

In one embodiment, as is illustrated in Fig. 2, at 30 document control server 12 receives a document request from the external interface for a first document. At 32, document control server 12 determines a user associated with the document request and authenticates the user. At 34, system 10 checks to see if the user is authorized to access the document requested. If so, at 36, system 10 retrieves the document from document server 14, cleans the document at 38 and, at 40, forwards the clean version of the document to the user. One embodiment of such a system and method is described next for a HyperText Transfer Protocol (HTTP) system.

When an HTTP or HTTPS connection request comes into document control server 12 there are three critical functions that must take place prior to returning the requested Web page: authentication, authorization, and internal connection. If either of the first two functions fail, the internal connection is not made. Then, once the internal connection has been made, document control server 12 must parse and "clean" the Web page prior to returning it to the requesting user.

Authentication

Authentication is fairly straight-forward and is of course visible to the end user. Following the HTTP protocol, when a user first enters a Uniform Resource Locator (URL) from their browser and the request is received at 30 (see Fig. 2) at server 12, a check is made at 32 for authentication. If basic authentication is being used, a check is made for authentication information in the HTTP header. If no username and password are found, the server returns a 401 error to the browser, telling the browser it needs to authenticate. The browser then pops up the box prompting the user to enter a username and password. When the HTTP request comes into the server, document control server 12 parses the username and password, comparing against its internal database of users and if it finds a match, lets control proceed onto the authorization step. If no match is found, document control server 12 returns an error code back to the Web server it is running under (e.g., Internet Information

Server or Netscape Enterprise Server) and the server will once again send back a 401 requesting the username and password. This process can happen 3 times and then the server will deny access. In one embodiment this authentication check is checked against a data base of known users as opposed to letting the
5 Web Server check against a user database it may have.

Authorization

Once document control server 12 has authenticated the request, it must, at 34, determine if the user is authorized to get to the URL they have requested. This authorization will fail if the URL the user requested is not in the list of
10 "allowed" URLs associated with the user. In one embodiment, each user is assigned one or more roles. Each role has access to a set of allowed URLs associated with that role.

In one embodiment each user has one or more roles associated with their user ID. For instance, they could be in the Marketing role, as well as the
15 Engineering role. In one such embodiment, each role is directly associated with an internal server; you can only define one role for each server. This means you could not have the Marketing role and the Engineering role going to the same physical internal server. Such an approach can simplify system design.

In another embodiment, more than one roles may be assigned to each
20 internal server. For example, a manufacturer may have all his reseller information on one server. One role, however, contains international resellers and another role contains domestic resellers. In such an embodiment, it would be advantageous to be able to define different sets of URLs on a single document server 14 that would allow for the different roles.

25 To complete the authorization portion, document control server 12 scans the list of allowed URLs for each role the user is in until it finds a match. If no match is made, an error condition is returned to the Web server indicating access is denied and the Web server in turn sends the appropriate error back to the browser.

30 In one embodiment, document control server 12 translates the URL prior to doing its search for a match on the URL. When an external business partner (user) enters the URL, they enter a URL where the first part of the URL points

document control server 12 and the second portion is the 'role' associated with that URL

e.g., https://www.<server ID>.com/Engineering/Standards/http_protocol.html

5

where

https = Secure http connection using SSL

www.<server ID>.com = DNS name of document control server 12 (this name is unique to the customer installation)

10 /Engineering = role associated with this particular URL

/Standards/http_protocol.html = actual web page on the internal web server

If the real intranet web server associated with the Engineering role were
15 engineer.abcd.com, then the translated URL that document control server 12 would search for is:

engineer.abcd.com/ Standards/http_protocol.html

Intranet Connection

If both the authentication and authorization phases completed successfully,
20 document control server 12 will open a TCP connection to the appropriate intranet server (engineer.abcd.com from the above example). Once the TCP connection has been made, document control server 12 generates an http request for the specific web page. The intranet server locates and returns the requested web page to document control server 12.

25 Parsing the Page

The pages returned by the intranet are categorized as either text or non-text. Examples of the latter are graphics, such as GIF or JPEG documents, sound objects, or executable objects, such as Java applets. Non-text pages are not parsed. They are forwarded back to the client browser unchanged. Text documents, such as HTML
30 formatted pages, however, contain embedded links that may need to be translated into their external equivalent. Embedded links fall into 3 groups, some of which require

translation, while others don't: relative path links, server path links and absolute path links.

Relative path links, which are of the form *subdir/page.html*, don't require translation because the browser will prepend the path based on the referrer's page. For
5 example, if the referrer's page was at:

http://www.document_control_server.com/Engineering/Standards/http_protocol.html

and the relative link was *ssl_protocol.html*, then the browser would prepend
10

http://www.document_control_server.com/Engineering/Standards/

to the link.

Server path links take the form of */Specification/wheel.html* and require
15 translation. This type of link points to a page that resides on the same server as the referrer's page, but with an absolute path starting at the root directory of the server. Assuming the same referrer's page as in the paragraph above, the translated link would be */Engineering/Specification/wheel.html*. (Note that the access string *http://* is not required because the browser will fill that in.) The translation is performed by
20 prepending the alias associated with the referrer's page, *Engineering*, in this case, to the path of the embedded link.

Absolute path links are full URLs, such as

<http://engineer.abcd.com/Performance/testdrive.html>
25

and require translation only if they point to a server that is in document control server 12's Alias table. The example link will get translated because it points to the server *engineer.abcd.com* that exists in document control server 12's Alias table as *Engineering*. The translation is done by replacing the intranet server's name by the
30 document control server 12's server name, followed by the alias of the intranet server. In this example, the translated URL would be

http://www.<server 12 ID>.com/Engineering/Performance/testdrive.html.

Links that point to pages on servers unknown to document control server 12 are not translated because they may well point to valid external sites, such as Yahoo, which should be left untouched. In one embodiment, therefore, these links are not translated. (Note that if the referrer's page came in through the Secure Socket Layer (SSL), i.e., the URL starts with *https://*, then the translated links will also have *https://*.)

On the other hand, such links could pose a security threat. That is, the link could be pointing to an intranet server that contains sensitive information, whose existence should not be revealed to external users. To counter this, in one embodiment document control server 12 includes a list of links which should be hidden from the outside world. Links found on such a list would be translated to something innocuous.

Redirection

When a page has moved, an intranet server may send a redirect status back to document control server 12. This means that document control server 12 has to translate the redirected address, similar to how embedded links are handled, before forwarding it to the client browser.

Architecture

A document access system such as system 10 illustrated in Fig. 1 enables users to grant outside organizations direct access to internal web data in a secure, simple and manageable way. It is essentially a secure window through which outside partners can view internal web data. If, as is shown in Figure 4, external interface 16 includes a firewall 40, system 10 also provides authenticated, authorized and view-customized business partner access to key Intranet servers via a standard web browser. Such a system enables users to easily, but accountably, grant authenticated partner access to internal web data, with complete control and authorization. Outside partners need only access a predefined URL in order to access an internal web page.

In one embodiment, as is shown in Figure 4, document control server 12 is installed inside firewall 40. In another embodiment, as is shown in Figure 5, document control server 12 is installed on a third network. Either way document control server 12 authenticates the outside user and then routes the request to a hidden, internal URL. The entire process is transparent to the outside user, and easily defined by the internal

document control server 12 user. This allows business partners direct access to the data, eliminating the time lag, redundancy, lack of integrity and the bottleneck within the MIS Dept. (Please note that if document control server 12 is installed inside the firewall as is shown in Figure 4, firewall 40 must be configured to restrict HTTP requests from any
5 external source so that they can only get through to server 12. Similarly, if document control server 12 is installed on a third network (as a type of demilitarized zone (DMZ)) as is shown in Figure 5, firewall 40 must be configured to restrict HTTP requests into internal network 20 so that they can only come from server 12.)

In a third embodiment, such as is shown in Figure 6, document control server 12
10 is installed outside firewall 40 and is accessed through the Secure Sockets Layer (SSL). Such an embodiment should be set up so that firewall 40 allows HTTP traffic *only* from document control server 12 into internal network 20.

To further reduce any bottleneck, in one embodiment document control server 12 includes the option for the actual "data owners" themselves to define which partners
15 have access to selective internal data. A data owner is a trusted individual within the organization that is empowered to grant Business Partners access privileges to Web pages on document servers 16. In one such embodiment, a Data Owner is assigned to one or more "roles," where a "role" represents the mapping alias assigned to one or the servers 16. A Data Owner can only add Business Partners or map URLs for the server
20 "role" to which the Data Owner is assigned.

For example, an employee working in the Accounting department would be assigned to an Accounting role (server). The Accounting Data Owner is only able to access the internal servers specified by the administrator. This prevents the Accounting Data Owner from mapping URLs on any other server such as the Marketing or
25 Engineering servers.

Once a Data Owner has been assigned to a role, he or she is able to perform the following tasks:

- Add, modify, or delete a Business Partner from that particular role
- Establish a user ID and password for a Business Partner for basic
30 authentication
- Post or map an internal URL for access by a Business Partner
- Delete URLs from a posted Go List

Delegation of such tasks to the data owners frees up MIS while also delegating data administration to those who understand the information best. In such an embodiment, the system administrator also defines general authentication rules and the list of eligible document servers 16.

- 5 Business Partners are somewhat-trusted end users. They can be granted controlled access to selected Web page structures on internal Web servers(s) such as document servers 16 once they have been provided with the following information:
- A URL connecting them to document control server 12.
 - A user ID and password to authenticate them to document control server
 - 10 12.
 - The name of the "menu tag" that they will select when they connect to document control server 12 that will retrieve the internal Web pages as specified by the Data Owner.

It is the Data Owners' responsibility to create and maintain a listing of Business Partners
15 that require access to the intranet servers they control and provide the Business Partner with the information they will need to access the selected intranet server(s).

Every Business Partner defined by a Data Owner is part of a "group." The "group" a Business Partner belongs to is directly related to the role a Data Owner has been assigned and what internal servers are associated with that role. These groups
20 control what URLs they are able to access on the internal servers.

A Business Partner can be assigned to multiple groups. For example, a Business Partner may belong to both the Marketing and the Sales groups. Data Owners manage their Business Partner accounts through the Business Partner list. From the Business Partner List, a Data Owner can establish a new Business Partner and modify or delete an
25 existing Business Partner to any groups that they control.

In one embodiment, a Business Partner List is accessed by clicking on a BP List button on a Data Owner Administration utility window.

In one embodiment, document control server 12 includes a go list processor 22 and a document processor 24 (see Fig. 3). Go list processor 22 determines if the user
30 has authorization to access said first document. Document processor 24 reads a document from document server 14, cleans the document as detailed above and forwards

a clean version of the document to the user. Go list processor 22 and document processor 24 will be discussed next.

a) The Go List

The Go list is used by the document control server 12 to determine which URLs
5 an authenticated Business Partner may be allowed to display. The Go List is unique to each role. It is identified by the rolename.data within the roles/ directory. In one embodiment, the Go List is managed by MIS. Such an embodiment does not, however, take advantage of the flexibility provided by the architecture of the present invention. Instead, it can be advantageous to permit individual data owners to determine the URLs
10 to be included in each Go List. Such an embodiment will be discussed next. In this example, documents are made available by the Data Owner and can be accessed by a user termed a "business partner (BP)".

In one such embodiment, the Go list contains data formatted as follows:

15 Real_url; MENU="menu_name"

where the Real URL is the actual URL (without the http://) used by document control server 12 to access that particular directory or file. The MENU parameter is always present. There can be a value within the quotes, or it can be empty. If there is a value
20 within the quotes, then document control server 12 will parse that value up and set up a link to that particular URL with the title of the link being the Menu Name - if the Business Partner goes to its menu page after it logs in.

Only allowed URLs are present within the go list. No other URLs are included.

Also, the Go List will permit the Business Partner to access any of the files under
25 a certain directory. For the time being, this is done by default on any URL that the user allows that ends with a "/" - to allow the Business Partner to access anything within the subdirectory - this is entered in the go list with the traditional '*' following the trailing slash. In one embodiment, the directory URL is kept intact and the Data Owner is given the option to cut and paste the path that the Data Owner wants the whole directory
30 included in. In such an embodiment Data Owners append the * to the directory name if they want everything within that directory accessible to the Business Partners within that role. In another embodiment, explicit "disallows" could be included to handle

documents the Data Owner wants to except from inclusion in the list of accessible documents.

b) The Mapping Code on Document Control Server 12

The next portion of the whole mapping design is where a lot of the real work comes into play. There is some code within document control server 12 that gets called when the user (Data Owner) wants to map a particular server to the Go List. In one embodiment, a graphical user interface is used to select URLs and business partners. In one such embodiment, this code gets activated by the Data Owner performing one of the following tasks:

- 10 * (A) Bringing up the Server Go List for the first time
- * (B) Clicking on a Node within the Go List Mapping Tree that has not yet been expanded and may have some children node
- * (C) Clicking on the Remap Button

At this point the GUI communicates to Server 12 via a Get URL request. The
15 Get URL request:

- (A) indicates to the server to load the GO list for a particular role
- (B) checks to see if the node has not already been expanded and that the node exists on this server (the front part of the URL indicating the server name is consistent) and that the node is of an html type (or directory type) - if the node matches all of these
20 criteria, then the GUI will indicate to the server to expand the particular URL for the particular role and
- (C) If the DO selects the Remap button, they are prompted to see if they want to remap that portion of the server down. If the DO selects yes, then a request to Remap with the currently selected URL and the role is sent to the server.

25 Server 12 then acts upon the request that it receives from the GUI

- (A) if the request was to load the GO list, then the server portion checks for the existence of a role_map.data file in the roles directory. If this file does exist, then all that is done is that file is sent to the GUI line by line as is. If the file does not exist, then the file is created and the mapping function is called. The mapping function is called
30 with the file pointer, the name of the URL (the server name) to be mapped, and a depth indicator of 1. (NOTE: This embodiment includes an option to go down multiple layers, however due to timeout issues it may be better to just go down one layer at a time and

let the user build the map as they see fit. If it goes down multiple layers than the mapping function must have the skills and capabilities to prevent the same URL from being mapped multiple times (the recursive nature of links and web spiders). The mapping code and its behavior is described below these ordered steps.)

5 Once the mapping code is done, then the URLs with their appropriate line syntax have been input and saved into the mapping file. The mapping file is then sent line by line to the GUI.

(B) If the request was to expand a node, the server code then calls the mapping function for the particular URL to be expanded. It opens up a temp file to be used to
10 write the information into. It then calls the mapping code with a file pointer to this temp file, the URL to be mapped, and a depth of 1. (NOTE: same code called as in condition A, just different parameters). Once the mapping code returns, the file is communicated line by line to the GUI and then the file is removed from the system.

(C) If the request was to Remap a particular server or URL, then things get a
15 little tricky. If the DO had chosen to remap the entire server, then the URL sent is the base URL - otherwise the URL sent was the URL that the DO wanted to remap from that point down. The same code is used regardless of the situation - just a different URL value. What happens is:

The current map file is copied over to the new map file until the line with the
20 URL to be remapped is read in.

At this point this line is parsed to determine the depth in the tree (root level is 0).

A remap function is called which then does the following (note, that this is complicated by the fact that the tree could be at varying depths with files having been
25 added or deleted and we would like to keep the prior shape and values of the tree when applicable)

Create a temp file to contain the intermediate results of the mapping

Call the map function with a pointer to the temp file, the URL, and a depth of 1.

It then compares the current map file line by line with the temp file.

30 If the URL at the current depth is not found in the respective depth in the new temp file, that URL and any URLs immediately following it with a depth greater than the current depth are removed (that initial file is missing)

If the URL at the current depth is found in the temp file, any URL lines in between the URL being searched for and the prior URL are new files and are added prior to the current URL in the map file. Their syntax lines indicate the current depth and default of disallowing that URL. Then the existing current URL line is left as is in the map file.

At this point the next URL in the map file needs to be examined, in examining this URL the URL line syntax is examined. The depth is the issue of primary concern. If the depth is the same as the current depth, then continue with this loop. If the depth is a depth deeper than the current depth, then this URL is a child of the prior URL and the prior URL then also needs to be remapped - the remap function is then called recursively with the prior URL. If the depth is less than the current depth, then we are no longer examining URLs that needed to be remapped and this function then returns.

After the final remap function returns, the remainder of the values are not to be touched and so they are copied over to the new map file as is.

After the server is done remapping, it then sends the whole newly mapped tree back to the GUI line by line. (NOTE: the server also then recreates the Go list to reflect the new values, information on creating the Go list from the mapping information is below.)

Finally, the GUI reads in the lines of information it gets from the server.

For (A) an initially loaded server, the GUI will create a tree examining each line of input. This is described in section 3 of this summary.

For (B) an expanded node, the GUI will create children nodes immediately below the expanded nodes by setting the depth according to the new tree and parsing each line of input. The parsing of the lines of input is described in section 3 of this summary.

For (C) a remapped server, the GUI will delete the previous tree and re-create the new tree (similar to A).

The code that does the mapping is a set of code pieces that loads the URL, parses any HTML that is returned, and creates a chain of information regarding the links. It then searches however deep is desired on each of the links.

- it uses some standard libraries to aid in parsing and getting the URL and HTML

- when it encounters a link, it stores that information in memory. At that time it also attempts to determine what sort of file it is - currently we only distinguish between the following: HTML, sound, graphic, external, video. It attempts to determine this based off of the hints surrounding it based on the filename and the surrounding html.
- 5 - If the file is an HTML file and the mapping code has not yet reached the requested depth of mapping, the mapping code will then recursively try to bring up the HTML file and parse through its contents, etc.
- As it comes across a file and finishes parsing it - it creates a syntax line that the GUI is expecting and writes this line to the file that it is passed. The line is as follows:
 - 10 - URL (just a tag)
 - `http://real_file_url` (this is the URL to the file that the mapping code downloaded and parsed - this is the file that will be allowed or denied by the DO)
 - depth (an integer that is to indicate the current depth that this file is in the tree - the lines are listed such that the tree can be loaded via a depth-first sort of algorithm - it
 - 15 continues down the left hand side while the depth is getting bigger, adding children from left to right as the depth is the same, and going back up the tree as the depth becomes smaller)
 - filename (this is the name of the file for that URL - a * is used if the filename cannot be determined (like in the case of a directory or the server)
 - 20 - file type (this is a character which corresponds to the filetypes that were previously mentioned)
 - state of node (this is a character which indicates if this node was in a collapsed or expanded state when the tree was saved - this is used by the GUI only, the mapping code always defaults this to C)
 - 25 - allowed (this is a character which indicates if this URL is to be allowed or disallowed, the mapping code defaults this to disallowed)
 - status of attaining the link (this is the HTTP status code of trying to access this link
 - it could have a value of 200 which means it was accessed OK, 404 meaning
 - 30 that this URL was not found, or a 0 indicating that this was not accessed)

- already mapped (this is a one character flag used to indicate if this URL was already mapped and exists previously in the tree or not, this is most useful in a multi-depth mapping search)

Finally, the server has one more responsibility with respect to the mapping code.

5 The server must handle writing out the saved data from the GUI when the DO requests to save the data. At this point, the data that is posted from the GUI to the server is written out line by line into the map file again. After this is done, the server deletes the prior Go List and parses through the map file a line at a time determining if that line is allowed. If the line is allowed the real_url part of the Map line is added (minus the
10 http://) to the Go List. Next, the server checks the line to see if a "Menu" tag has been appended, if so, the server then adds the appropriate Menu tag to the Go list. Otherwise, it just adds a null menu tag to the Go list. As mentioned previously, currently if the real_url ends in a '/' an '*' is appended to the end of the real_url line to indicate that the user can access the entire directory - see Section 1 for further information.

15 After the Go List has been saved, the server is re-initialized with the new values - thus allowing immediate access or denial of access to the Business Partners for that role.

The Directory Map with the GUI

As mentioned previously, in one embodiment the GUI reads in and interprets the
20 given line in order to creating nodes to represent the mapped server as a tree structure to the DO. Once such tree structure can be seen in Figure 7.

The GUI communicates with server 12 as previously discussed and gets back well-defined lines of data. It parses the data and creates a node for each line provided by the Go list. It then looks at the expanded and collapsed feature to determine whether
25 that node should be expanded or collapsed. It also looks at the file type to associate an image with that node. This image should allow the user to better determine what sort of nodes they are giving the Business Partners access to. The image is also determined by the return status from the mapping process - if a status of 404 is specified, then that link is determined to be broken (at least from document control server 12) and a broken link
30 icon is displayed next to it. Finally, if that node is currently allowed, then the green ball is displayed next to it - to given the illusion of "green light means Go" to the DO.

The directory tree is created in a depth first fashion. It reads in each line examining the current depth. If the depth of the new node to be inserted is greater than the current depth than the node is inserted as a child of the previously inserted node. If the depth is the same, than it is inserted as a sibling. If the depth is less, then the tree is
5 parsed back until the depth of the tree is at the same value of the depth and the node is inserted as a sibling at that level. The server is considered to be at the root level and thus has a depth of 0.

The text value that shows up with the node in the tree is the real_url value minus the http:// appended by the Menu Tag and Value if there is one for that node. A node
10 may have a menu tag without being allowed.

The Data Owner (DO) can then traverse the tree examining the various links and determining what to allow and what to disallow. If they allow a directory - everything beneath that directory will be allowed. There currently is no mechanism for handling exceptions yet. If they allow files, then that file is allowed. If they have checked any
15 other "include on allow" options (currently we offer, include all Gifs, Audio, Video, HTML, All Links) - then the files immediately beneath (once again only one layer deep) are automatically turned to allow if they are of the corresponding type. One note is that external files will never be "allowed" - as they do not exist on the server and thus it does not make sense for the DO to be allowing or disallowing those files.

20 When a data owner selects a link, they have the option of specifying the menu tag to be shown. If they do not specify one, then it is left blank. If they do then it is assigned for that node only. In order for it to get assigned to that node, the DO will have to select Allow or Disallow.

When Data Owners have finished making their changes they can save or cancel
25 their mapping values. If they cancel then nothing that they did since their last save or remap will be saved. If they hit save, the values are communicated back to the server and the map file and the Go List File are updated as described previously.

Installation

Prior to install: Before fully installing and configuring document control server
30 12, the customer must have a digital certificate (for SSL encryption or transmissions) as well as set up and configure a server such as Internet Information Server (on NT) or Netscape Enterprise Server (on UNIX Solaris).

Installations: MIS installs document control server 12 onto the IIS or NES server and configures the firewall to allow HTTP access to the document control server 12 server.

Definition of end users: Document control server 12 offers organizations the option to delegate administration to end users who control the actual data (Data Owners) rather than forcing more work onto MIS. The data owner's access is defined by the network administrator. The data owner then maps which servers can be accessed. This data is stored in the document control server 12 "go list." The program code implementing document control server 12 is now completely installed, and ready for use.

Define business partner access: At this point, in order for an outside partner to access data, he/she must be granted access by the data owner. The data owner simply accesses the document control server 12 "data owner" GUI via a standard Java-enabled web browser. He/She can then define the new partner via role-based administration or explicitly choose which URLs may be accessed.

Outside users will not have access to any internal URL that is not specifically listed, even if there are embedded links in a URL for which access was granted. However, users can define access to a particular URL and all sub-pages as well.

Future partner access: After one role has been defined, future partners need only be added to that role, rather than requiring a whole new access to be defined.

Business partner access: All the outside partner has to do is type in the defined URL with any standard web browser. The partner will then be prompted for a user ID and password. Once these are entered, the partner will see a list of accessible URLs.

25 Back End Databases

It is important to understand that document control server 12 simply passes HTML information. This means document control server 12 does not have a problem passing CGI scripts and other dynamic content. Where this becomes particularly confusing is the access and authentication to back end databases via an application gateway.

One of document control server 12's greatest values is to allow outsiders access to ever-changing information such as order processing, shipping, etc. Much of this data

is stored in large back-end databases with an application gateway on the front end. The application gateway puts an HTML front end on the database and allows Intranet users to query required information. Typically, a user only needs to enter a customer account number to access this information. However, in order to give outside users direct
5 access, many organizations need to require some level of authentication to this process.

When document control server 12 passes an outside partner to any Intranet URL, the user is authenticated as a unique document control server 12 user, however, that user ID is not passed on to the Intranet server. Therefore, direct access to back end databases cannot be defined for each document control server 12 Business Partner. It will be
10 necessary to create an HTML-based sign-in screen. This is a simple process and offers an opportunity for resellers and professional services to add value to the product sale.

Many Internet web servers handle restricted access in slightly different ways. If the user is not known, the web server responds with a 401 error. The user's browser then displays a standard screen requesting user ID and password. The user types these in
15 and is granted access.

It is important to understand that document control server 12 cannot process this transaction. For security reasons, only HTTP traffic can pass through document control server 12. Any authentication must be HTTP based, as mentioned above.

In one embodiment, document control server 12 is installed on a standard web
20 server running IIS (NT) or NES (Solaris). It requires no changes to the current infrastructure, and no "agents" or "clients" to be installed on any web servers or browsers.

Administration and data owner usage is accessed via document control server 12's Java user interface. This allows access via any web browser that supports Java
25 (e.g., Internet Explorer 4.0, Netscape 4.0). Outside partner access is also accomplished via a standard web browser.

Operating with Third Party Firewalls

Document control server 12 can be used in conjunction with a firewall to add an additional layer of security to Business Partner communications via the Web. Two
30 components need to be considered when determining the location of document control server 12: Domain Name Service (DNS) and routing.

Depending on the deployment option preferred and the capabilities of firewall 40, there are up to four different methods for routing traffic to, and through, server 12:

1) Redirected proxy— For added security on external to internal connections, a redirected proxy can be configured on your firewall to redirect the inbound connection requests. When a Business Partner on the external network attempts to connect to document control server 12, firewall 40 intercepts the request and establishes a connection to server 12. This rerouted connection hides the actual destination from the Business Partner requesting the connection.

2) Transparent proxy— A transparent proxy can be set up through firewall 40 to document control server 12. From the Business Partners' perspective it will appear as though they are connecting directly to server 12 and not connecting to the firewall first.

3) Directly to document control server 12— If document control server 12 is installed on the external side of firewall 40 (as in Figure 6), connection requests will be routed directly to server 12. In such an embodiment, server 12 authenticates the Business Partner, and passes the request through firewall 40. Firewall 40 then retrieves the requested Web page(s) from the specified document server 16.

4) Through a third network— (some firewalls allow a "third network" capability, (sometimes called the DMZ or the Secure Server Network). The three deployment scenarios discussed above still apply in a "three network" environment; however, additional firewall configuration is necessary to ensure that the required name resolution (DNS), and routing are still possible.

Security Features

SSL encryption. In one embodiment, data transmitted between the partner and web server is SSL-encrypted to prevent a sniffer from gathering information from the connection.

Document control server 12 server encrypted. Data stored on document control server 12 server such as user IDs, the go list, and partner profiles are all encrypted to prevent unauthorized access.

Password and user ID authentication. In one embodiment, document control server 12 supports password and user IDs for authorization. Stronger encryption could also be used.

Granular Access Controls

Business partners can only access internal URLs to which explicit access is given. If an accessible URL has embedded links to pages to which explicit access has not been granted, the partner cannot connect to them. However, if an embedded link is
5 to an outside server, such as www.yahoo.com, in one embodiment access will not be restricted.

Internal URLs and IP Address Are Hidden

To ensure the security of the internal network and web pages, internal URLs and IP addresses are hidden from outside access. Partners type in a predefined URL and are
10 presented with a list of accessible internal URLs. When a link is selected from the list Document control server 12 then maps to the internal URL. The internal URL and IP address are never displayed for the partner to see.

System Requirements, Compatibility and Performance

Considerations for performance and reliability include amount of cache memory,
15 CPU power, BUS speed, amount of RAM, speed of memory chips, bus architecture (IDE, EIDE, PCI etc.), hard drive capacity, and hard drive quality (seek and access speeds). The following table identifies the basic characteristics of minimum, recommended and ideal server configurations to run document control server 12.

System Component	Minimum	Recommended	Ideal
CPU	Pentium 166	Pentium 200	Pentium Pro 200
RAM	32 MB	48 MB	64 MB
5 Hard disk	1 GB	2 GB	4 GB
Platform	IIS 3.0 (NT 4.0) or NES 3.0 (Solaris 2.5.1)		
Browser	Java enabled web browser MS Internet Explorer 4.0 or higher Netscape Navigator 4.0 or higher with Netscape's JDK 1.1 patch		
Other	CD-ROM, 3.5" diskette, Color monitor, Keyboard, Mouse		

10

Document control server 12 enables users to easily, but accountably, grant authenticated partner access to internal web data, with complete control and authorization. Outside partners need only access a predefined URL in order to access an internal web page.

15

The following examples provide a better idea of how document control server 12 can be used to meet a variety of needs.

Example – Manufacturing (Order Processing)

Manufacturing companies process thousands of orders every day. In order to keep up with competition and to achieve the highest levels of quality, customers/partners need to know the immediate status of an order to the minute. Many companies today have moved to just-in-time inventory systems to reduce overhead and costs. Document control server 12 can grant access directly to an order-processing page that connects directly into an order-processing database. The order-processing agents (data owners) can define what data customers/partners have direct access to. As a result, the customer knows immediately the status of an order. The supplier also saves money by eliminating the need to replicate data or take phone calls asking for updates.

25

Example – Distribution

A distribution environment operates an order-processing and shipping department very similar to manufacturing. However, distribution also requires various types of information to be distributed to different partners, such as pricing and quantity

30

breaks. Document control server 12 allows a company to customize the view each distributor or reseller sees, such as pricing or quantity breaks.

Example – Financial Services

Financial institutions process millions of transactions a day with a large number of outside partners. These include the purchase and sale of assets as well as order/sale confirmation, etc. Today, many of these transactions require a third party to set up a secure certificate. Document control server 12 can speed up this entire process by allowing an agent to immediately allow an outside customer or partner access to trading information in minutes, and without the need for third-party intervention.

10 Example – Health Services

Health care and insurance organizations process thousands of claims each day. Partners need a secure way to pass medical information and process it into a company's systems. For example, a doctor treats a patient who has Blue Cross/Blue Shield. That doctor needs to know if the patient's insurance covers the treatment, then process the claim after the treatment is given, and finally check on the status of payment once a claim is submitted. With document control server 12, Blue Cross/Blue Shield can give the doctor's office access to their internal list of insured patients, as well as the status of current claims. The company no longer needs to replicate this data to a DMZ or SSN Internet server or handle a phone call. The doctor's office can also securely fill out a web-based claim form over the Internet to process the claim for treatment.

20 Example – Government Agency

It is necessary for various government agencies and departments to frequently share sensitive data. One example is the CIA and various law enforcement agencies. The FBI, DEA, ATF and other agencies must routinely check into the files of various personnel and public citizens. Typically, this requires these agencies to send a paper request for information to the CIA. The CIA must then search for the relevant information and then send a copy back to the requesting agency.

With document control server 12, the FBI and other agencies can be given direct access to the CIA files that might be relevant such as histories and fingerprint analysis databases. This can save time and money.

Document control server 12 offers several advantages over current methods such as cost savings, improved customer service and leveraging of the current infrastructure.

Current methods for passing data to outside partners are expensive, slow and unreliable. Document control server 12 offers the information to partners faster, easier and cheaper. It also more tightly integrates partners, thus improving business relations. Document control server 12 also leverages the benefits of current technology such as the Internet
5 and Intranet.

Other business advantages of document control server 12 include: it reduces overhead and costs; it eliminates the need to copy content to a web server within the DMZ or external network; it offers spontaneous, dynamic user-managed content; it eliminates the wait for an IS manager to update data or post on a web server; it
10 eliminates integrity and replication issues; it more tightly integrates partners; and its open architecture allows access without the need to alter current technology.

Although specific embodiments have been illustrated and described herein, it will be appreciated by those of ordinary skill in the art that any arrangement which is calculated to achieve the same purpose may be substituted for the specific embodiment
15 shown. This application is intended to cover any adaptations or variations of the present invention. Therefore, it is intended that this invention be limited only by the claims and the equivalents thereof.

What is claimed is:

1. A method of limiting access from an external network to documents stored on an internal network, the method comprising the steps of:
 - building a client list, wherein the step of building a client list includes the step of
 - 5 assigning each client to a role;
 - building a document list naming documents available to clients assigned to the client's role;
 - receiving a request for a document stored on the internal network;
 - associating the request with a client;
 - 10 determining if the requested document is on the list of documents; and
 - if the requested document is on the list of documents, fetching the requested document as a proxy and sending the requested document to the client.
2. The method according to claim 1, wherein each document has a unique URL and
- 15 wherein the document list includes the URL of each available document, wherein the step of determining includes the step of determining if the URL of the requested document is included in the document list.
3. The method according to claim 2, wherein the step of building the document list
- 20 includes the step of displaying available documents in a tree structure within a graphical user interface, wherein the step of displaying includes the step of querying a document control server to obtain a current version of the document list.
4. The method according to claim 1, wherein the step of associating the request
- 25 with a client includes the step of authenticating the client.
5. The method according to claim 4, wherein each document has a unique URL and
- wherein the document list includes the URL of each available document, wherein the
- step of determining includes the step of determining if the URL of the requested
- 30 document is included in the document list.

6. The method according to claim 5, wherein the step of building the document list includes the step of displaying available documents in a tree structure within a graphical user interface, wherein the step of displaying includes the step of querying a document control server to obtain a current version of the document list.

5

7. A document control system, including:

an internal network;

an external interface;

a document server connected to the internal network, wherein the document

10 server controls access to a plurality of documents, including a first document; and

a document control server, wherein the document control server receives a

document request for the first document, determines a user associated with the

document request and authenticates the user, wherein the document control server

includes a go list processor for determining if the user has authorization to access said

15 first document and a document processor for reading the first document from the

document server, cleaning the first document and forwarding a clean version of said first document to the user.

8. The document control system according to claim 7, wherein the external

20 interface includes a firewall connected to an external network, wherein the document control server communicates to the document server through the firewall.

9. The document control system according to claim 8, wherein the document

control server is connected to the external network and communicates to the firewall

25 through the external network.

10. The document control system according to claim 8, wherein the document

control server is connected to a third network and communicates to the firewall through the third network.

30

11. The document control system according to claim 7, wherein the document

control server is connected to the internal network and wherein the external interface

includes a firewall connected to an external network, wherein the firewall is configured to receive the document request and to route the document request to the document control server.

- 5 12. The document control system according to claim 7, wherein the document processor acts as a proxy to hide access to the first document.
13. The document control system according to claim 7, wherein the external interface includes a telephone interface into which a business partner can dial to gain
10 access to the document control server.
14. The document control system according to claim 7, wherein the document control server includes means for translating links embedded in the first document.
- 15 15. The document control system according to claim 7, wherein each user is assigned one or more roles and wherein the go list processor restricts access to documents as a function of the role under which the user attempts to access the document.
- 20 16. A document control system, including:
an internal network;
an external interface;
a document server connected to the internal network, wherein the document server controls access to a plurality of documents, including a first document;
25 a document control server; and
a data owner interface for building a document list of available documents;
wherein the document control server receives a document request from the external interface for the first document, determines a user associated with the document request and authenticates the user; and
30 wherein the document control server includes a go list processor for determining, based on the document list, if the user has authorization to access said first document.

17. The document control system according to claim 16, wherein the data owner interface includes a graphical user interface which displays the document list in a tree structure, wherein the graphical user interface queries the document control server to obtain a current version of the document list.

5

18. The document control system according to claim 16, wherein the document control server further includes a document processor for reading the first document from the document server, cleaning the first document and forwarding a clean version of said first document to the user.

10

19. The document control system according to claim 18, wherein the document processor acts as a proxy to hide access to the first document.

20. The document control system according to claim 19, wherein the data owner interface includes a graphical user interface which displays the document list in a tree structure, wherein the graphical user interface queries the document control server to obtain a current version of the document list.

21. The document control system according to claim 16, wherein the external interface includes a firewall connected to an external network.

22. The document control system according to claim 16, wherein the external interface includes a telephone interface into which a business partner can dial to gain access to the document control server.

25

23. In a system having an internal network and an interface to an external network, a method of handling requests from the external network for documents stored on the internal network, the method comprising:

- defining one or more users;
- 30 defining documents accessible to the users;
- receiving a document request from the external network;
- determining a user associated with the document request;

authenticating the user associated with the document request;
determining if the user associated with the document request has permission to
access the document requested;
if the user associated with the document request has permission to access the
5 document requested, retrieving the document requested from the internal network.
cleaning the document of embedded links and delivering the document to the user
associated with the document request.

24. The method according to claim 23, wherein each document has a unique URL
10 and wherein determining if the user associated with the document request has
permission to access the document requested includes:

accessing a document list listing the URL of each available document; and
generating an error message if the document requested is not on the document
list.

15

25. The method according to claim 23 or 24, wherein defining documents accessible
to the users includes assigning each user to one or more roles and limiting access to
documents as a function of role and wherein determining if the user associated with the
document request has permission to access the document requested includes determining
20 if users in the role associated with the document request have permission to access the
document requested.

26. The method according to claim 23, wherein each document request includes an
HTTP header and wherein authenticating the user associated with the document request
25 includes retrieving authentication information from the HTTP header.

27. The method according to claim 23, wherein cleaning the document of embedded
links includes looking for a server path link and replacing the server path link with a link
to an alias.

30

28. The method according to claim 23, wherein cleaning the document of embedded
links includes looking for an absolute path link, determining if the absolute path link is a

link which should be hidden and, if the absolute path link is a link which should be hidden, replacing the absolute path link with a different link.

29. In a system having an internal network and an interface to an external network, a method of handling requests from the external network for documents stored on the internal network, the method comprising:

5 defining a plurality of users, including a first and a second user;
assigning each user to one or more roles, wherein assigning includes assigning the first user to a first role and the second user to a second role;
10 defining documents accessible to the users, wherein defining includes limiting access to documents as a function of the roles assigned to the user;
receiving a document request from the external network;
determining a user and a role associated with the document request;
authenticating the user associated with the document request;
15 determining if users in the role associated with the document request have permission to access the document requested;
if users in the role associated with the document request have permission to access the document requested, retrieving the document requested from the internal network and delivering the document to the user associated with the document request.

20

30. The method according to claim 29, wherein each document has a unique URL and wherein determining if the user associated with the document request has permission to access the document requested includes:

accessing a document list listing the URL of each available document; and
25 generating an error message if the document requested is not on the document list.

31. The method according to claim 29, wherein retrieving the document requested includes cleaning the document of embedded links.

30

NOT FURNISHED UPON FILING

1/6




FIG. 1




FIG. 2




FIG. 3




FIG. 4

4/6




FIG. 5




FIG. 6

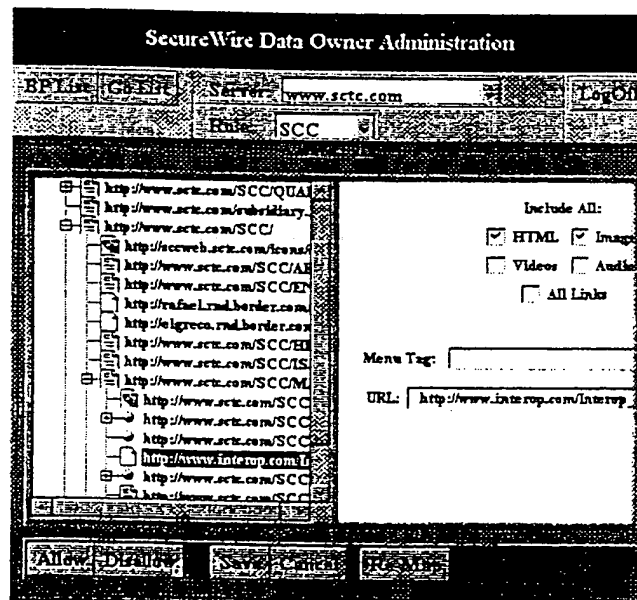


FIG. 7

INTERNATIONAL SEARCH REPORT

Int. l. Application No

PCT/US 99/03382

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04L29/06 G06F1/00

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 96 13113 A (SECURE COMPUTING CORP) 2 May 1996 see abstract; figures 2-5 see page 5, paragraph 2 - paragraph 3 see page 20, paragraph 3 - page 21, paragraph 3	1-31
A	KAHAN J: "A CAPABILITY-BASED AUTHORIZATION MODEL FOR THE WORLD-WIDE WEB" COMPUTER NETWORKS AND ISDN SYSTEMS, vol. 27, no. 6, 1 April 1995, pages 1055-1064, XP002039534 see the whole document	1-31

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"Z" document member of the same patent family

Date of the actual completion of the international search.

14 June 1999

Date of mailing of the international search report

21/06/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 apo nl,
Fax: (+31-70) 340-3016

Authorized officer

Powell, D

INTERNATIONAL SEARCH REPORT

International Application No.

PCT/US 99/03382

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	VINTER S T: "EXTENDED DISCRETIONARY ACCESS CONTROLS" PROCEEDINGS OF THE SYMPOSIUM ON SECURITY AND PRIVACY, OAKLAND, CALIFORNIA, APRIL 18 - 21, 1988, 18 April 1988, pages 39-49, XP000673340 INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS	
A	EP 0 697 662 A (IBM) 21 February 1996	
A	EP 0 811 939 A (WEBTV NETWORKS INC) 10 December 1997	

INTERNATIONAL SEARCH REPORT

Information on patent family members

Inter. Application No

PCT/US 99/03382

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9613113 A	02-05-1996	US 5864683 A AU 3888595 A EP 0787397 A	26-01-1999 15-05-1996 06-08-1997
EP 0697662 A	21-02-1996	CA 2154020 A JP 8087454 A	16-02-1996 02-04-1996
EP 0811939 A	10-12-1997	AU 3375197 A JP 10228437 A WO 9746943 A	05-01-1998 25-08-1998 11-12-1997

THIS PAGE BLANK (USPTO)